

## UC Research Cyber-infrastructure: The Business Case for a Pilot

The urgency of creating a world-class UC Research Cyberinfrastructure<sup>1</sup> (CI) is reflected in a rapidly escalating need for greater compute capability and capacity to do competitive research and the escalating energy costs and space needs associated with these resources. Investments by peer institutions are impacting the competitiveness of our campuses in attracting the best students and staff and in garnering extramural support. Many UC faculty and the UC Vice Chancellors for Research (VCRs) and Chief Information Officers (CIOs) strongly support a UC Research CI. UC has already implemented a nationally recognized grid infrastructure that supports remote collaboration, access, and management of compute resources regardless of geographical location.

In the absence of shared UC research computing facilities, there has been a proliferation of distributed computer cluster systems throughout the campuses in sub-optimal, remodeled locations. These are extremely costly in terms of facilities, power, cooling, and space, costing 3 to 4 times more than efficient data center space. More importantly, these resources are not leveraged with each other to generate greater research capability and are often underutilized. If we continue with “business as usual”, these costs will continue to spiral upward and imperil our research enterprise. We currently estimate that clusters distributed throughout the campuses occupy a footprint of over 25,000 sq.ft., displacing scarce instructional and office space across UC. Collectively, distributed clusters already exceed the computing power at major supercomputer centers. Industry research indicates that cluster computing is growing at approximately 30% per year. At UC we project an increase in computing power by 2015 that will be the equivalent of adding at least a second major supercomputer. Additionally, UC can expect to add as much as 80,000 sq.ft. in new data center space at a cost of \$160 - \$320 million. Over half of these resources will be needed for research computing and storage. Power costs, even with efficient data center designs, are expected to grow by \$10 – 15 million in annual costs. Five UC campuses will reach a crisis point in data center facilities and will need new space within two years.

In sharp contrast to our persistent investment in distributed resources, many UC competitors are now investing heavily in common computational resources for their researchers. Institutions, like U. Illinois, U. Texas, U. Southern California, and U. New Mexico each recently made investments of \$5-15M to centralize their resources into a large capability and capitalize upon the resulting economies of scale. UC will simply not be competitive without investing in research computing.

To address our pressing CI needs, the UC VCRs and CIOs have collaborated on a plan for staged, coordinated investment that we believe will build centralized, world-class CI capabilities and at the same time manage escalating costs with much greater cost-benefit effectiveness. The CI pilot began with a call for proposals in April, 2008 for computationally-intensive projects that would 1) advance UC research in priority areas, such as global health and environmental science; 2) make the projects more competitive for obtaining extramural support; and 3) nucleate new communities of cyber-enabled researchers in areas like the social sciences, arts, and humanities. Over 30 proposals were received from almost all campuses and national laboratories, and about two dozen of these are projects that are suitable for inclusion in the pilot.

The preliminary estimate for funding the CI pilot is \$5 million for hardware, networking and storage in existing data center space provided by the Lawrence Berkeley Laboratory and the San Diego Supercomputer Center. Annual recurring costs for staff support will be \$1.5 million. These estimates will be refined based on specifications developed collaboratively with the research groups. In aggregate, the pilot projects should be better positioned to compete for extramural support. While this first phase of the pilot is focused on computation, succeeding stages will address data storage, visualization and virtualization.

We believe UC is, right now, within the optimum window in time for investing in this pilot CI project. By doing so, UC will create the cyber-infrastructure and partnerships that will yield substantial dividends for the cyber-intensive research environment that lies ahead. By consolidating compute and data center resources, UC can provide greater research capability and at the same time significantly lower facility costs. We respectfully urge UC’s support of this critically important initiative.

---

<sup>1</sup> CI is the coordinated integration of a diversity of hardware, software and human resources that includes network (campus, system, state and national), computational clusters, collaborative clusters, shared storage systems, co-location facilities, collaboration / virtualization tools, system administration support, code librarians, coding experts, grid access tools, visualization tools and resources and licensed software.